UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF GEORGIA
ATLANTA DIVISION

MOBILE NETWORKING SOLUTIONS, LLC, Plaintiff, vs. E*TRADE FINANCIAL HOLDINGS, LLC, Defendant.	**CASE NO. 1:21-cv-_____** **JURY TRIAL DEMANDED**

COMPLAINT FOR PATENT INFRINGEMENT

Mobile Networking Solutions, LLC ("MNS") files this Complaint for Patent Infringement against E*Trade Financial Holdings, LLC ("E*Trade") for infringement of U.S. Patents Nos. 7,543,177 and 7,958,388 relating to large-scale data storage, processing, and management.

PARTIES

1. MNS is a limited liability company organized and existing under the laws of the State of Texas with its principal place of business at 1400 Preston Road, Suite 483, Plano, Texas 75093.

2. E*Trade is a Delaware limited liability company headquartered in Arlington, Virginia, and registered to do business in Georgia.

3. E*Trade may be served through its registered agent, Corporation Service Company, 40 Technology Pkwy South, #300, Norcross, Georgia 30092.

4. E*Trade is an online financial services firm that provides a trading platform geared for digitally inclined investors and traders. E*Trade executes millions of financial transactions a day generating large amounts of data.

5. To manage and access this vast amount of data, E*Trade has developed a cloud data lake to provide enterprise-wide reporting and analytic capabilities across platforms.

6. E*Trade stores data in Hadoop clusters using HDFS. *See* https://www.elastic.co/elasticon/tour/2018/chicago/elastic-at-etrade.

JURISDICTION AND VENUE

7. This is an action under the patent laws of the United States, 35 U.S.C. §§ 1, *et seq*. and namely §§ 271, 281, and 284-285, for infringement by E*Trade of U.S. Patent Nos. 7,543,177 (the "'177 Patent") and 7,958,388 (the "'388 Patent") (collectively, the "Patents-in-Suit").

8. This Court has subject-matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

9. E*Trade is subject to general and specific personal jurisdiction of this Court based upon its regularly conducted business in this district giving rise to this action and has established minimum contacts with this forum such that the exercise of jurisdiction over E*Trade would not offend traditional notions of fair play and substantial justice.

10. E*Trade, directly and through subsidiary business units, has committed and continue to commit acts of infringement in this district pursuant to 35 U.S.C. § 271(a) by making, using, selling, offering to sell, testing, deploying, and exercising control and obtaining beneficial use in this district of products and services that infringe the asserted MNS patents.

11. Venue is proper in this judicial district pursuant to 28 U.S.C. § 1400(b) and 28 U.S.C. § 1391.

12. Venue is proper in this district pursuant to 28 U.S.C. § 1400(b) as E*Trade maintains an office and data center in this district at 1650 Union Hill Road, Alpharetta, GA 30005.

THE MNS PATENTS

13. MNS is the owner by assignment of all right, title, and interest in and to U.S. Patent Nos. 7,543,177 and 7,958,388 (the "Asserted Patents"), both titled, "Methods and Systems for a Storage System."

14. A true and correct copy of the '177 patent is attached as Exhibit A.

15. A true and correct copy of the '388 Patent is attached as Exhibit B.

16. MNS possesses all rights of recovery under the Asserted Patents.

17. The Asserted Patents issued from continuations of Application No. 10/284,199 filed on October 31, 2002.

18. The U.S. Patent Office issued the '177 Patent on June 2, 2009, after a full examination based upon an application filed by inventors Melvin James Bullen, Steven Louis Dodd, William Thomas Lynch, and David James Herbison.

19. The Examiner stated the following reasons for allowing the claimed subject matter of the '177 Patent:

> Regarding claim 1, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, a switch controller that executes software, including a routing algorithm and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm.

> Regarding claim 26, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, a management system determining a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface and the management system removing from service the memory section from which the fault message was received by changing the routing algorithm.
>
> Regarding claim 40, the prior art does not disclose or reasonably suggest, in combination with the remaining limitations, programmable means for switching data being transmitted between the means for storing and one or more interfaces based on a routing algorithm and means for receiving the fault message, removing from service the means for storing from which the fault message was received by changing the routing algorithm.

20. The U.S. Patent Office issued the '388 Patent on June 7, 2011, after a full examination based upon an application filed by the same inventors.

21. The Examiner stated the following reasons for allowing the claimed subject matter of the '388 Patent: "the prior art does not teach or reasonably suggest providing, by the management system, the routing algorithm to the switch controller and determining, by the management system in response to the detecting, a new routing algorithm that redirects data for the memory device to a replacement memory device; and providing the new routing algorithm to the switch controller."

22. The Abstract of the Asserted Patents describes the claimed subject matter as being directed to "[a] storage system that may include one or more memory sections, one or more switches, and a management system . . . [t]he memory sections include memory devices and a section controller capable of detecting faults with the memory section and transmitting messages to the management system regarding

detected faults. The storage system may include a management system capable of receiving fault messages from the section controllers and removing from service the faulty memory sections . . . [a]dditionally, the management system may determine routing algorithms for the one or more switches."

23. Figure 6 in the specification of the Asserted Patents is a functional diagram exemplifying the claimed subject matter:

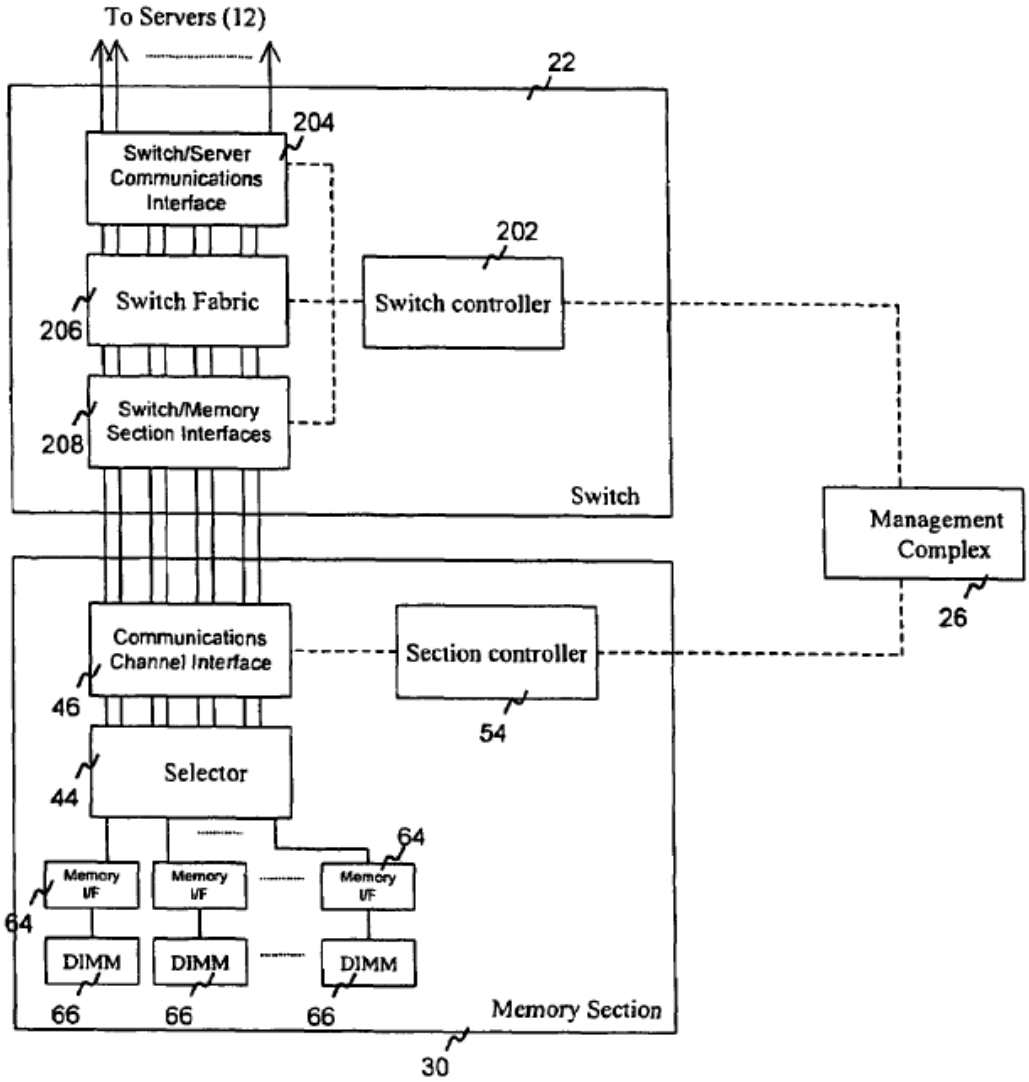

24. The inventors recognized and noted in the specification that large-scale

storage systems suffered from problems in throughput for high-volume, real-time applications.

25. In operation, the switches, memory sections, and management system of the Asserted Patents receive fault messages from the memory section controllers and remove from service the memory section from which the fault message was received, and the management system may further determine an algorithm for use by a switch fabric in interconnecting the memory sections and external device interfaces and instruct the switch to executed the determined algorithm. '177 Patent at 2:21-34.

26. Those of skill in the art at the time of the inventions claimed in the Asserted Patents would recognize that the claimed subject matter addresses performance limitations inherent in disk storage technologies such as input/output bottlenecks and improves network operations in the event of signal and/or equipment failure by improving fault management in storage systems.

27. A person of ordinary skill in the art would recognize that the claimed subject matter of the '177 and '388 Patents is directed to improvements to large-scale storage management systems, in particular large capacity storage required to deliver high-volume, real-time application memory access. *See* '177 at 2:3-7.

28. For example, online transaction processing (OLTP) systems facilitate and manage transaction-oriented applications, typically for data entry and retrieval transaction processing in which the system responds immediately to user requests. Common examples of OLTP systems are time-critical financial transaction processing systems, eTrading, brokerage, electronic banking (e.g., ATM transactions), and e-commerce systems.

29. To a person of ordinary skill in the art at the time of the invention,

transaction processing means information processing that is divided into individual, indivisible operations called transactions. Each transaction must succeed or fail as a complete unit. A transaction cannot remain in an intermediate state.

30. A disadvantage, recognized by the inventors, of prior art systems is their susceptibility to service disruptions as a consequence of their online, real-time functionality. Operations can be severely impacted if the system is unavailable due to data corruption or system failures.

31. In the '177 and '388 Patents, "fault management" refers to attempting to detect faults and take corrective action in response to the detection of a fault. '177 at 8:13-22.

32. The consequence of service disruptions is heightened in large-scale storage systems comprising mechanical disk drives and solid-state disk storage systems that may be combined in cached disk arrays under common program control. *See* '177 at 1:31-54. These systems may house an organization's mission-critical application data, so throughput and fault resilience are particularly important.

33. A person of ordinary skill in the art at the time of the invention would understand that throughput in the context of the claimed inventions refers generally to workload, the volume of transactions per time period that the storage system is capable of handling. Effective fault management is important for maintaining throughput and ensuring system operations is not disrupted by events such as, for example, a catastrophic failure of a memory section. '177 at 7:48-53.

34. Examples of memory section faults include a failure to write to a magnetic storage device due to a power failure ('177 at 8:23-29), failure of a memory device in a cached disk array ('177 at 8:40-49), or an entire memory section fault ('177 at 8:50-63).

35. Before the inventions claimed in the '177 and '388 Patents, large-scale storage systems for real-time applications lacked effective fault detection and management for providing and maintaining sufficient throughput and fault resiliency.

The Asserted Claims

36. Claim 1 of the '177 Patent is directed to a storage system apparatus featuring, among other structural requirements and functionality, memory sections, switches including a switch controller that stores and executes a routing algorithm configuring a switch fabric, and a management system capable of: receiving fault messages; inactivating memory sections corresponding to a fault message by changing the routing algorithm; and providing the determined routing algorithm to the switch controller for execution.

37. Claim 1 of the '177 Patent recites:

1. A storage system, comprising:

(a) one or more memory sections, including:

(i) one or more memory devices having storage locations for storing data, and

(ii) a memory section controller capable of detecting faults in the memory section and transmitting a fault message in response to the detected faults;

(b) one or more switches, including:

(i) one or more interfaces for connecting to one or more external devices;

(ii) a switch controller that executes software, including a routing algorithm; and

(iii) a selectively configurable switch fabric connected to one or more memory sections and the one or more interfaces and interconnecting the memory sections and the one or more interfaces based on the routing algorithm stored in the switch controller; and

(c) a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm, and wherein the management system is further capable of determining and changing the routing algorithm for use by the selectively configurable switch fabric in interconnecting the memory sections and the one or more interfaces, providing the determined routing algorithm to the switch controller, and instructing the switch controller to execute the determined routing algorithm.

38. A person of ordinary skill in the art at the time of the invention would recognize that claim 1 of the '177 Patent is directed to an improved large-scale storage management system for delivering high-volume, real-time application memory access and fault management.

39. Claim 1 represents an advance over the prior art by, among other things, providing a memory section controller capable of detecting faults and transmitting a fault message in response to detected faults, switches including a switch controller

and selectively configurable switch fabric connecting memory sections based upon a routing algorithm stored in the switch controller, and a fault management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the received fault message by changing the routing algorithm. '177 at 28:42-67.

40. '177 Patent Fig. 6 is a functional diagram illustrating the novel system of claim 1:

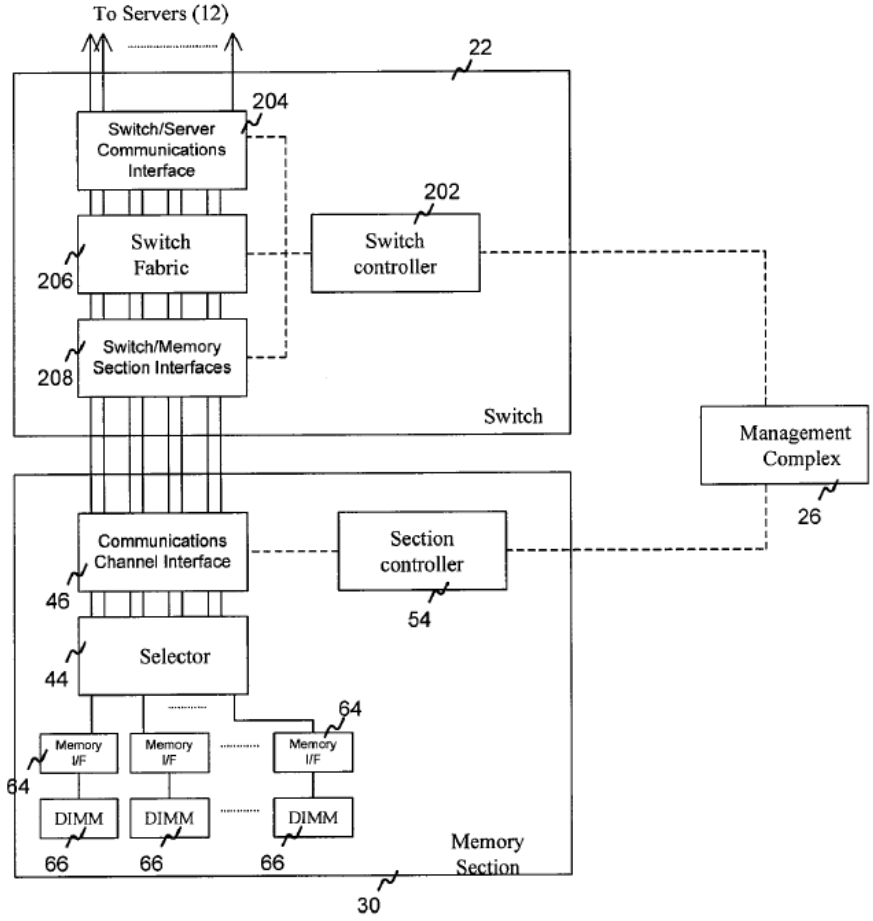

Fig. 6

41. Conventional switches may be implemented in a switch fabric of known type (e.g., "IP switch fabric, an FDDI switch fabric, an ATM switch fabric, and Ethernet switch fabric, an OC-x type switch fabric, or a Fibre channel switch fabric," '177 at 13:24-29); however, they are arranged, managed, and operated in an unconventional and novel way according to '177 claim 1.

42. In claim 1, a novel switch controller executes a routing algorithm promulgated by the management system, which is capable of determining and changing the routing algorithm and instructing the switch controller to execute the determined routing algorithm to accommodate and overcome faults in a memory section.

43. In the event the section controller detects a fault (e.g., through the use of the Header/Test interface 63, referring to '177 Patent Fig. 5), the controller may transmit fault information such as time, component, and type of fault to the management system. '177 at 12:54-58.

44. A person of ordinary skill in the art would understand that the section controller and switch controller of claim 1 are unconventional and each operates together with the components of the claimed system in an unconventional way to achieve improved throughput.

45. The claimed subject matter of claim 1 includes a selectively

configurable switch fabric connected to memory sections and interconnected interfaces based upon a routing algorithm. A switch fabric is the physical interconnection architecture that directs data from an incoming interface to an outgoing interface. '177 at 5:63-65.

46. In the '177 and '388 Patents, the term "configuration" encompasses the various possible operating states of each component of the storage system. '177 at 7:39-42. Operating states refer to operational states defined by parameters such as, for example, how often a component (e.g., memory section) sends performance statistics, the list of events that cause a component (e.g., memory section) to report a fault alarm and the type of alarm reported. '177 at 7:42-53.

47. The selectively configurable switch fabric of claim 1 may be changed in response a detected fault. For example, a failed memory section is removed from the fabric, and data is written to a different memory section to ensure throughput remains above operational thresholds. Configuring the switch fabric in response to fault detection marks an improvement to storage systems that existed before the inventions claimed in the '177 and '388 Patents.

48. Claim 1 of the '177 Patent includes a novel management system capable of receiving fault messages from the memory section controllers and removing from service the memory section from which the fault message was

received and determining an algorithm for use by the switch fabric and execution by switch controllers to maintain operation of the system.

49. A person of ordinary skill in the art would recognize the claimed management system as a novel component of the claimed storage system having fault-management functionality that was not known at the time of the invention. She would also understand the fault-management functionality delivered by the management system delivers the benefits and advantages described in the patents as a marked improvement to storage system throughput and reliability.

50. For example, the management system is capable of performing fault management by recognizing an operational failure of a memory section or part of a memory section and re-mapping data to working memory sections. '177 at 8:12-29.

51. Examples of bad block re-mapping functionality provided by the novel management system are detailed in the specification ('177 at 8:30-49):

> For example, if the control processors 34 discover that block 65,000 in memory section 30-2 does not read correctly, the control processor 34 may decide to remap block 65,000 in memory section 30-2 to block location 1,999,998 in memory section 30-2. The control processor 34 may then direct the CDA 16 to read the data block and cause it to be written in location 1,999,998 in memory section 30-2. Once completed, the control processor 34 may inform the switches 22 and memory section 30-2 that block 65,000 may now be read from location 1,999,998.

> As another example of bad block remapping, if for example only one memory device on a memory section is faulty, a control processor 34 in the management complex 26 may inform the section controller 54 about the bad device, determine where the data on the faulty memory device is backed-up (e.g., CDA 16), and direct the backed-up data to be loaded into a replacement memory device on the same memory section or on a different memory section. In the latter case, the management complex also informs the switch about the data being relocated to a new memory section.

52. The management system of claim 1 did not exist in storage systems before the inventions claimed in the ’177 and ’388 Patents.

53. Inclusion of the management system of claim 1 in the claimed storage system architecture was unconventional at the time of the invention and materially changes the operation of the other system components including the switches, switch controllers, switch fabric, memory section controllers, and memory sections.

54. Conventional and well-known storage systems at the time of the inventions claimed in the ’177 and ’388 Patents did not manage routing algorithms in a management system distinct from a switch controller to centralize configuration of a switch fabric.

55. A person of ordinary skill in the art would understand and recognize the improvements delivered in the claimed system over the prior art systems. The capability of the management system to change routing algorithms to direct transactions to “healthy” memory sections and inactivate memory sections

corresponding to the fault message delivers substantial advantages in operational throughput and reliability of large-scale storage systems.

56. Advantages delivered by the claimed inventions include the management system's capability of "bringing new memory sections into service and taking memory sections out of service independently of other functions that the management complex performs and without materially affecting the operation of other memory sections [] or adversely affecting the overall performance of the storage hub." '177 at 7:64-8:2. One of ordinary skill in the art would recognize that this functionality is particularly advantageous in the event that "a memory section [] has failed, then the faulty memory section [] may be replaced and a new one brought into service." *Id*. at 8:7-10.

57. The centralized fault management procedures claimed in the MNS patents provide concrete improvements in operation. "[I]f for example, only one memory device on a memory section is faulty, a control processor [] in the management complex [] may inform the section controller [] about the bad device, determine where the data on the faulty memory device is backed-up [], and direct the backed-up data to be loaded into a replacement memory device on the same memory section or on a different memory section." *Id*. at 8:40-47. This functionality achieves the inventors' objective "to provide sufficient throughput for high-volume,

real-time applications." *Id*. at 2:3-5.

58. Claim 13 of the '177 Patent is directed, generally speaking, to a method of using a novel storage system by storing data in locations in a memory device in a memory section, determining a routing algorithm for use by a switch controller to configure a selectively configurable switch and providing the routing algorithm to a switch controller, and detecting a fault in regard to data stored in the memory device and transmitting a fault message to the management system, which in turn removes the relevant memory section from service by changing the routing algorithm.

59. The process and series of steps recited in claim 13 was not conventional or well-known at the time of the inventions claimed in the '177 Patent.

60. Operating a storage system as recited in claim 13 was not possible before the invention of the '177 Patent because the system components used in the method did not exist and using conventional, known storage system components available at the time of the invention would not result in performance of the claimed method.

61. Claim 13 of the '177 patent recites:

> 13. A method for use in a storage system, comprising:
>
> (a) storing data in a storage locations in a memory device, the memory device included in a memory section;

(b) a management system determining a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface;

(c) the management system providing the determined routing algorithm to the switch controller and instructing the switch controller to execute the determined routing algorithm;

(d) the selectively configurable switch connecting the memory section to the interface based on the routing algorithm;

(e) detecting by a memory section controller a fault in regard to the data stored in the memory device and transmitting a fault message in response to the detected fault to the management system; and

(f) receiving the fault message at the management system; and the management system removing from service the memory section from which the fault message was received by changing the routing algorithm.

62. In the method of claim 13, a novel management system determines a routing algorithm executed by a switch controller to configure a selectively configurable switch in connecting the memory section where data is stored and an interface.

63. Performance of the method of claim 13 configures a switch connecting a memory section and interface.

64. A person of ordinary skill in the art would understand the method of claim 13 is substantially different than the making or use of the apparatus of claim 1 at least in part due to the recited detection by a memory section controller of a fault in regard to the data stored in the memory device, which is not recited in claim 1.

65. Removing a faulty memory section from service, which is performed by a management system by changing the routing algorithm, was not a routine or conventional process at the time of the inventions claimed in the '177 Patent because management systems to the extent they existed at the time were incapable of performing the claimed step.

66. A person of ordinary skill in the art would understand that claim 13 of the '177 Patent describes a process of using a storage system that delivers advantages and operational benefits that was neither convention, well-known nor understood at the time of the inventions as improving throughput and reliability of large-scale storage systems.

67. Claim 1 of the '388 Patent is directed to a storage system apparatus featuring, among other structural requirements and functionality, memory sections, switches including a switch controller that executes a routing algorithm configuring a switch fabric, and a management system capable of: receiving fault messages; inactivating memory sections corresponding to a fault message by changing the

routing algorithm; and providing the determined routing algorithm to the switch controller.

68. Claim 1 of the ’388 Patent is materially different than claim 1 of the ’177 Patent. “In this embodiment, the switch controller 202 and memory section interfaces 208 need not be included in the switch 22, and the management complex 26 of the storage hub 10 exercises direct control over the switch fabric 206 and server interfaces 204.” ’388 at 13:50-59.

69. ’388 Patent Fig. 7 illustrates the functional diagram of the system of claim 1.

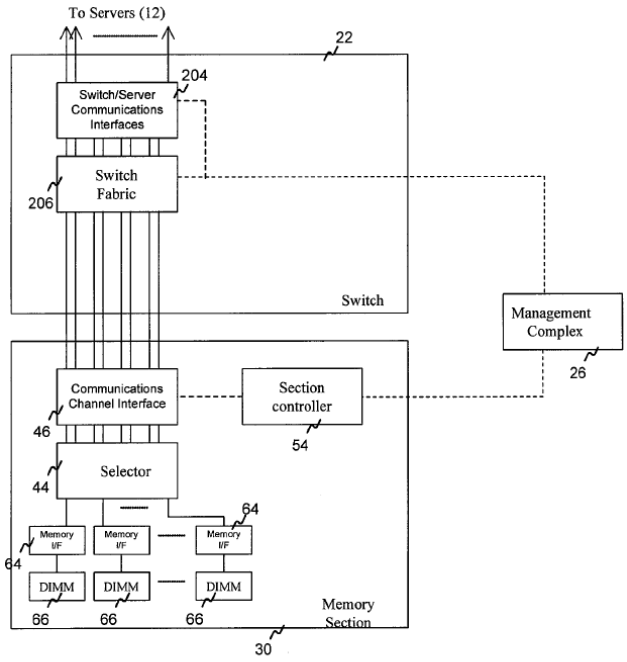

Fig. 7

70. A person of ordinary skill in the art would recognize that claim 1 of the

'388 Patent is directed to an improved large-scale storage management system for delivering high-volume, real-time application memory access and fault management.

71. ’388 Patent claim 1 represents an advance over the prior art by, among other things, providing a memory section controller capable of detecting faults and transmitting a fault message in response to detected faults, switches including a switch controller and selectively configurable switch fabric connecting memory sections based upon a routing algorithm, and a fault management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the received fault message by changing the routing algorithm.

72. Conventional switches may be implemented in a switch fabric of known type (e.g., “IP switch fabric, an FDDI switch fabric, an ATM switch fabric, and Ethernet switch fabric, an OC-x type switch fabric, or a Fibre channel switch fabric,” ’388 at 13:24-29); however, they are arranged, managed, and operated in an unconventional and novel way according to ’388 claim 1.

73. In ’388 Patent claim 1, a novel switch controller executes a routing algorithm promulgated by the management system, which is capable of determining and changing the routing algorithm and instructing the switch controller to execute

the determined routing algorithm to accommodate and overcome faults in a memory section.

74. In the event the section controller detects a fault (e.g., through the use of the Header/Test interface 63, referring to '388 Patent Fig. 5), the controller may transmit fault information such as time, component, and type of fault to the management system. '388 at 12:54-58.

75. A person of ordinary skill in the art would understand that the section controller and switch controller of claim 1 are unconventional and each operates together with the components of the claimed system in an unconventional way to achieve improved throughput.

76. The claimed subject matter of claim 1 includes a selectively configurable switch fabric connected to memory sections and interconnected interfaces based upon a routing algorithm. A switch fabric is the physical interconnection architecture that directs data from an incoming interface to an outgoing interface. '388 at 5:65-67.

77. In the context of the claims, configuring the switch fabric encompasses various operational states defined by parameters such as, for example, how often a component (e.g., memory section) sends performance statistics, the list of events that cause a component (e.g., memory section) to report a fault alarm and the type of

alarm reported. ’388 at 7:39-53.

78. The selectively configurable switch fabric of ’388 Patent claim 1 may be changed in response a detected fault. For example, a failed memory section is removed from the fabric, and data is written to a different memory section to ensure throughput remains above operational thresholds. Configuring the switch fabric in response to fault detection marks an improvement to storage systems that existed before the inventions claimed in the ’177 and ’388 Patents.

79. Claim 1 of the ’388 Patent includes a novel management system capable of receiving fault messages from the memory section controllers and removing from service the memory section from which the fault message was received and determining an algorithm provided to a switch controller to configure the switch fabric.

80. A person of ordinary skill in the art would recognize the claimed management system as a novel component of the claimed storage system having fault-management functionality that was not known at the time of the invention. She would also understand the fault-management functionality delivered by the management system delivers the benefits and advantages described in the patents as a marked improvement to storage system throughput and reliability.

81. The management system of claim 1 is a novel structural component of

the claimed storage system and did not exist in storage systems before the inventions claimed in the ’388 Patents.

82. Inclusion of the management system of claim 1 of the ’388 Patent in the claimed storage system architecture was unconventional at the time of the invention and materially changes the operation of the other system components including the switches, switch controllers, switch fabric, section controllers, and memory sections.

83. Before the ’388 Patent invention, storage systems did not manage routing algorithms in a management system distinct from a switch controller to centralize configuration of a switch fabric.

84. In the embodiment claimed in ’388 Patent claim 1, the communication channel interface 46 (*see* Fig. 7, above) of the memory section 30 directly connects to the switch fabric 206. This arrangement of the system components was not conventional or well-known at the time of the inventions, and a person of ordinary skill in the art would recognize this embodiment as novel at the time.

85. A person of ordinary skill in the art would understand and recognize the improvements delivered in the system claimed in ’388 Patent claim 1 over the prior art systems. The capability of the management system to change routing algorithms to direct transactions to “healthy” memory sections and inactivate

memory sections corresponding to the fault message delivers substantial advantages in operational throughput and reliability of large-scale storage systems

86. Claim 2 of the ’388 Patent is directed to a method for use in a storage system reciting steps that include storing data in locations in a memory device in a memory section, determining a routing algorithm executed by a switch controller that is provided by the management system to configure a configurable switch connecting the memory section to an interface, detecting a fault associated with the data in the storage locations in the memory device, and determining by the management system in response to the detecting of a fault a new routing algorithm that redirects data to a replacement memory device.

87. The process and series of steps recited in claim 2 was not conventional or well-known at the time of the inventions claimed in the ’388 Patent.

88. Performing the method of claim 2 was not possible before the invention of the ’388 Patent because the system components used in the method did not exist and using conventional, known storage system components available at the time of the invention would not result in performance of the claimed method.

89. Claim 2 of the ’388 patent recites:

2. A method for use in a storage system, comprising:

(a) storing data in storage locations in a memory device, the memory

device included in a memory section;

(b) determining, by a management system, a routing algorithm for use by a switch controller that executes software, including the routing algorithm;

(c) providing, by the management system, the routing algorithm to the switch controller;

(d) executing, by the switch controller, the routing algorithm, to configure a configurable switch connecting the memory section to an interface;

(e) detecting a fault associated with the data in the storage locations in the memory device

(f) determining, by the management system in response to the detecting, a new routing algorithm that redirects data for the memory device to a replacement memory device; and

(g) providing the new routing algorithm to the switch controller.

90. In the method of claim 2, a novel management system determines a routing algorithm executed by a switch controller to configure a selectively configurable switch.

91. Performance of the method of claim 2 configures a switch connecting a memory section and interface.

92. A person of ordinary skill in the art would understand the method of claim 2 is substantially different than the making or use of the apparatus of claim 1

or the method of '177 claim 13 at least in part due to the recited determining of a routing algorithm by a management system and a new routing algorithm redirecting data to a replacement memory device in response to detection of a fault associated with the data in the storage locations in the memory device, which is not recited in claim 1 of the '388 Patent or claim 13 of the '177 Patent.

93. Determining and providing a new routing algorithm by the management system in response to a detected fault was not a routine or conventional process at the time of the inventions claimed in the '388 Patent because management systems to the extent they existed at the time were incapable of performing the claimed step.

94. Modifying a routing algorithm by a management system and providing a new routing algorithm to a switch controller was neither a routine, conventional nor well-known method at the time of the invention of the subject matter claimed in '388 Patent claim 2. Conventional systems at the time of the invention were incapable of operating as recited in claim 2.

95. A person of ordinary skill in the art would understand that claim 2 describes a process of fault detection and management that delivers advantages and operational benefits that was neither well-known nor understood at the time of the inventions as improving throughput and reliability of large-scale storage systems.

96. The claimed subject matter of the Asserted Patents is particularly applicable to improve the operation of parallel processing technologies in big-data distributed storage systems such as the Hadoop Distributed File System (HDFS).

Hadoop Distributed File System

97. The Hadoop Distributed File System (HDFS) is used for storage and processing of large data files across a cluster of storage hardware.

98. HDFS is a distributed, reliable, scalable and highly fault-tolerant file system for data storage, and E*Trade uses HDFS for long-term storage of user and content data that E*Trade analyzes to gain knowledge of user preferences and behavior.

99. An HDFS instance may consist of hundreds or even thousands of servers (DataNodes) that each store part of a large data file.

100. HDFS features high fault tolerance and automatic fault recovery making it suitable for deployment on commodity hardware and particularly valuable to E*Trade for reliable access to data.

101. Operational advantages of HDFS include efficient processing by executing application instructions near the subject data. HDFS's cluster design and input/output pathing minimizes network congestion and increases throughput.

102. HDFS handles big data, typically 10-100GB or more with diverse data

types including structured and unstructured data, economically distributing the computational load across multiple DataNodes.

103. HDFS DataNodes are a cluster of computers capable of executing the workload components such as storing HDFS data blocks and performing block replication.

104. Distributing the computing load across DataNodes requires multiple servers having access to the data, and HDFS meets this need by ensuring that the entire calculation process does not terminate when an error occurs within a HDFS cluster.

105. The NameNode is responsible for keeping track of file system metadata including a list of blocks in an HDFS file and a list of DataNodes.

106. MapReduce decomposes the processing task of a large data set query for processing on multiple running nodes. As a result, E*Trade is able to analyze massive datastores. For example, E*Trade is able to query terabytes of data for specific actions.

107. E*Trade's Hadoop/HDFS deployment ensures data availability because it is self-healing and fault tolerant.

108. In the event of a fault (i.e., a lost DataNode), the NameNode consults metadata, finds affected data, consults a Rack Awareness script, and instructs the DataNode to replicate. This HDFS process is described pictorially below:

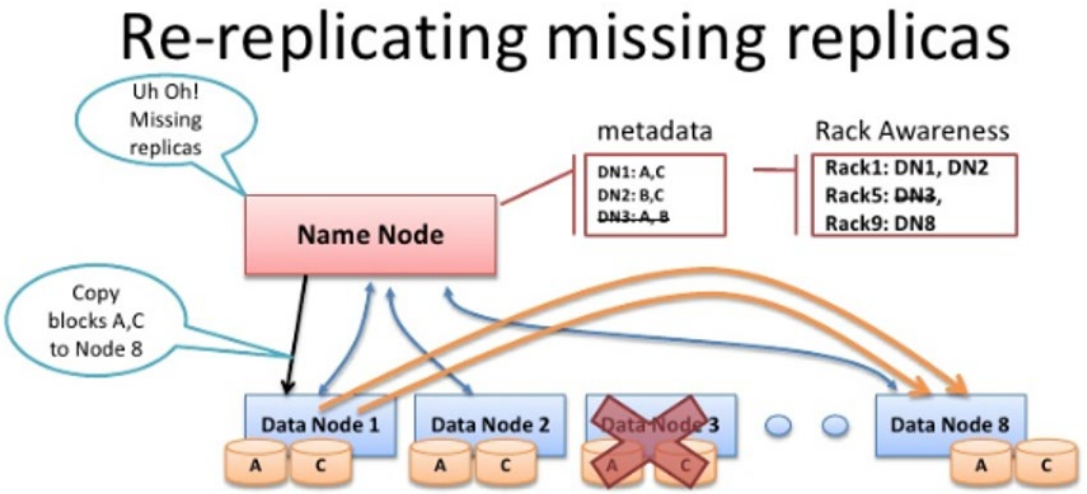

109. In E*Trade's implementation of HDFS, the NameNode manages the file system namespace and regulates access to files by clients and DataNodes manage storage attached to the nodes they run on.

COUNT I

INFRINGEMENT OF U.S. PATENT NO. 7,543,177

110. MNS re-alleges and incorporates by reference the preceding paragraphs as if stated here.

111. E*Trade has and continues to infringe at least claims 1 and 13 of the '177 Patent.

112. E*Trade makes, uses, sells, offers for sale, and/or imports the E*Trade

data lake using HDFS (the "Accused Instrumentalities").

113. The Accused Instrumentalities embody and practice the subject matter claimed in the asserted claims of the '177 Patent.

114. Asserted claim 1 of the '177 Patent recites a storage system, comprising: one or more memory sections, including: one or more memory devices having storage locations for storing data, and a memory section controller capable of detecting faults in the memory section and transmitting a fault message in response to the detected faults; one or more switches, including: one or more interfaces for connecting to one or more external devices; a switch controller that executes software, including a routing algorithm; and a selectively configurable switch fabric connected to one or more memory sections and the one or more interfaces and interconnecting the memory sections and the one or more interfaces based on the routing algorithm stored in the switch controller; and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm, and wherein the management system is further capable of determining and changing the routing algorithm for use by the selectively configurable switch fabric in interconnecting the memory sections and the one or more interfaces, providing the determined routing algorithm to the switch controller,

and instructing the switch controller to execute the determined routing algorithm.

115. Asserted claim 13 of the '177 Patent recites a method for use in a storage system, comprising: storing data in a storage locations in a memory device, the memory device included in a memory section; a management system determining a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface; the management system providing the determined routing algorithm to the switch controller and instructing the switch controller to execute the determined routing algorithm; the selectively configurable switch connecting the memory section to the interface based on the routing algorithm; detecting by a memory section controller a fault in regard to the data stored in the memory device and transmitting a fault message in response to the detected fault to the management system; receiving the fault message at the management system; and the management system removing from service the memory section from which the fault message was received by changing the routing algorithm.

116. The Accused Instrumentalities, and HDFS implementations on E*Trade data lake, are storage systems.

117. A typical architecture of a Hadoop cluster features Slave nodes for

storage and the NameNode that oversees and coordinates the data storage function.

118. In normal operation, the Accused Instrumentalities implementing HDFS store data blocks in a DataNode's (memory section) local file system that uses storage including memory devices (e.g., HDD, SSD). The memory devices store data in physical storage locations (e.g., HDD sectors, SSD blocks).

119. The Accused Instrumentalities include a management system that determines a routing algorithm for use by a switch controller that executes software, including the routing algorithm, to configure a selectively configurable switch in connecting the memory section and an interface.

120. In normal operation, the Accused Instrumentalities implementing HDFS manage the HDFS NameSpace (e.g., by operation of the HDFS NameNode daemon) and map data file names to sets of data blocks, map data blocks to specific DataNodes, and map DataNodes to specific racks in the HDFS cluster.

121. In the Accused Instrumentalities, NameNode NameSpace tables and resultant NameNode instructions based on them (i.e. the I/O path a HDFS client uses to read/write a specific data block) are routing algorithms used by the HDFS NameNode (switch controller) that controls how specific HDFS I/O requests traverse the HDFS cluster.

122. Consistent with the asserted claims, the Accused Instrumentalities

implementing HDFS achieve high fault tolerance by ensuring persistence of file system metadata.

123. In the accused HDFS instances, the HDFS namespace is stored by the NameNode, which uses a transaction log called the EditLog to persistently record every change that occurs to file system metadata.

124. For example, creating a new HDFS file in the Accused Instrumentalities causes the NameNode to insert a record into the EditLog. Changing the replication factor of a file also causes a new record to be inserted into the EditLog. The NameNode stores the EditLog, and the entire file system NameSpace, including the mappings and system properties, is stored by the NameNode.

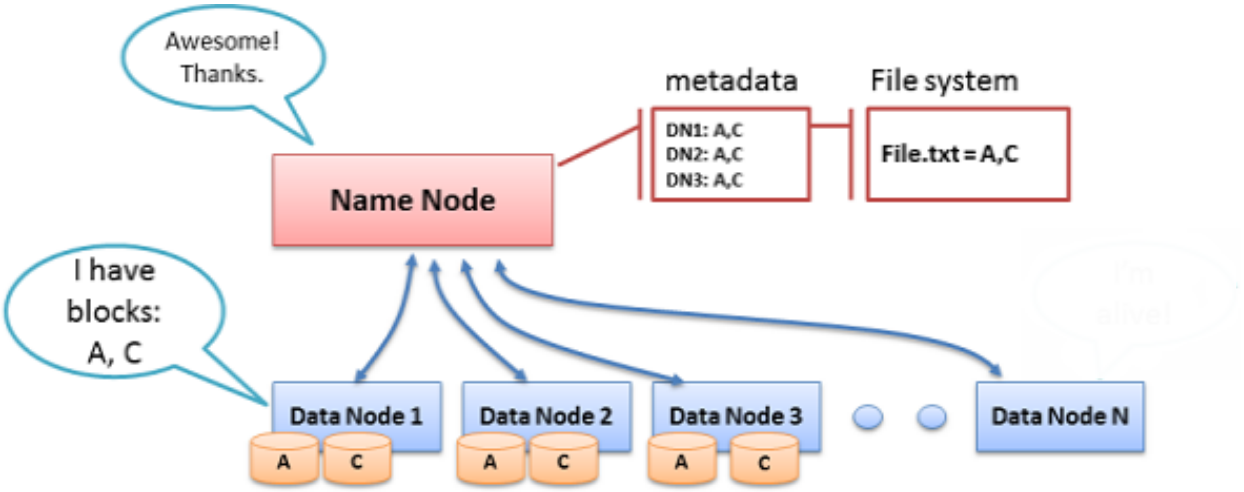

125. The figure above provides a representative diagram of the switch controller (NameNode), routing algorithm (metadata and file system), memory section (DataNode N), and memory devices (devices labeled A, C).

126. In the accused HDFS implementations, the NameNode daemon determines the routing algorithm by processing the metadata tables in response to HDFS client Read and Write operations (exemplified in the figures below).

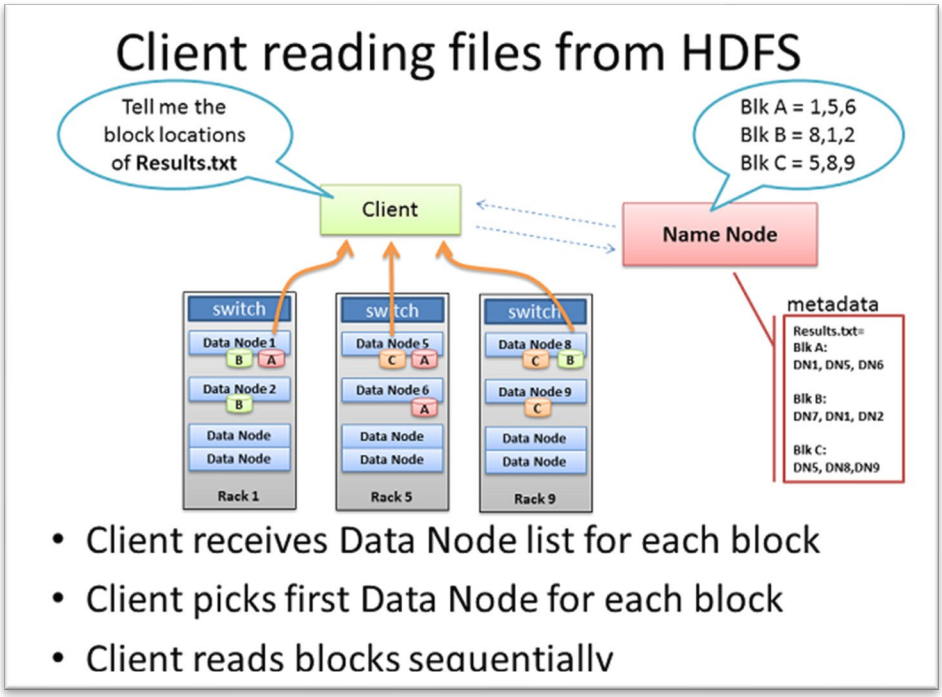

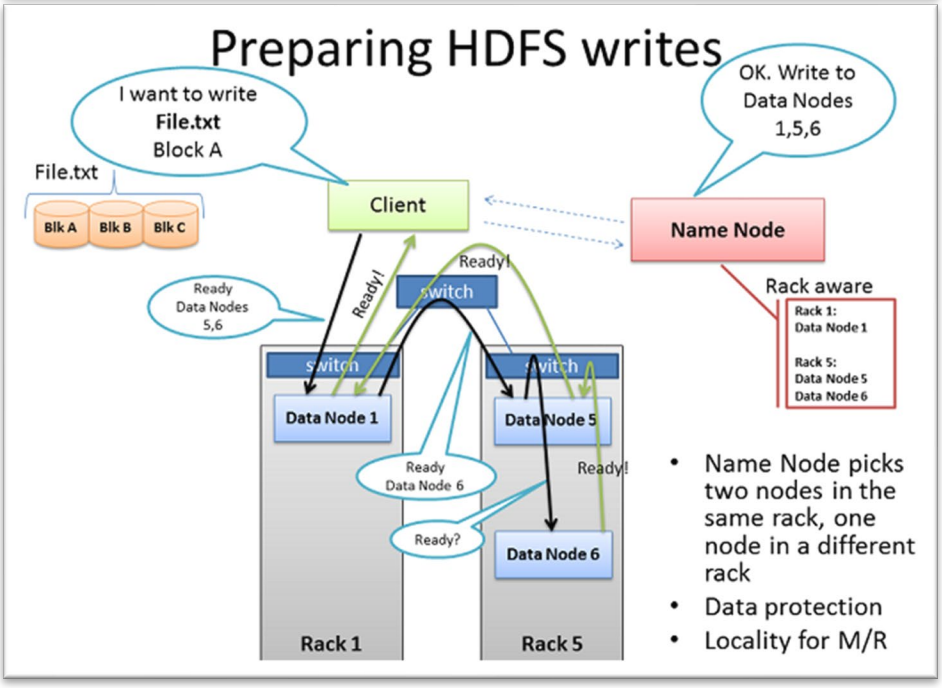

127. The Accused Instrumentalities include network switches.

128. In large clusters, the Accused Instrumentalities spread the nodes across multiple racks. Nodes of a rack share a switch, and these rack switches, which are selectively configurable, are in turn connected by one or more core switches.

129. In the Accused Instrumentalities, selectively configurable rack switches connect HDFS data nodes (memory sections) to an interface.

130. In the event of an HDFS I/O request, the rack switch routes the request to the proper HDFS data node in accordance with the HDFS file system NameSpace that includes the mapping of blocks to files.

131. In normal operation of the Accused Instrumentalities, a memory section controller (e.g., data node daemon) detects a fault in regard to data stored in the memory device and a fault message is transmitted to the management system (e.g., HDFS NameNode) in response to the detected fault.

132. During normal operation, each DataNode periodically sends a heartbeat message to the NameNode. If a subset of DataNodes lose connectivity with the NameNode, the NameNode detects the fault by the absence of a heartbeat message and marks the affected DataNodes as dead and ceases forwarding any new I/O requests to them. The NameNode tracks which blocks need to be replicated due to a fault and initiates replication when necessary.

133. By default, the heartbeat is transmitted every three seconds, set by dfs.heartbeat.interval.

134. In addition to detecting a fault by monitoring heartbeats, HDFS DataNodes create threads that run a DataBlockScanner object that scans the data blocks (and replicas) stored in the DataNode to detect faults.

135. The Name Node daemon receives the fault message in the NameNode (management system) due to either a disruption in heartbeats from a DataNode or receipt of a DataBlockScanner report indicating a fault.

136. During normal operation of the Accused Instrumentalities, upon detecting a dead DataNode (e.g., a DataNode with no heartbeat) the NameNode daemon (management system) bypasses the dead DataNode and instead sends I/O requests to the other DataNodes storing replicas of blocks that were stored on the dead DataNode. If a corrupted block is detected (e.g., via DataBlockScanner) the NameNode daemon (management system) marks the block replica as corrupt and then schedules a copy of the block to be replicated on another DataNode, which results in an updated HDFS NameSpace (a new routing algorithm) so its replication factor is back at the expected level. Thus, during normal operation, the management system removes from service the memory section from which a fault message was received by changing the routing algorithm.

137. E*Trade is on notice of the infringing products, services, features, and how E*Trade operates the Accused Instrumentalities to perform the claimed methods and use the claimed apparatuses.

138. E*Trade's infringing conduct has damaged MNS.

139. E*Trade is liable to MNS in an amount that adequately compensates it for Defendants' infringement, which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT II

INFRINGEMENT OF U.S. PATENT NO. 7,958,388

140. MNS re-alleges and incorporates by reference the preceding paragraphs as if stated here.

141. E*Trade has infringed and continues to infringe at least claims 1 and 2 of the '388 Patent by making, using, selling, and/or offering to sell the Accused Instrumentalities.

142. The Accused Instrumentalities embody and practice the asserted claims of the '388 Patent.

143. Asserted claim 1 of the '388 Patent recites a storage system, comprising: one or more memory sections, including one or more memory devices having storage locations for storing data, and a memory section controller capable

of detecting faults in the memory section and transmitting a fault message in response to the detected faults; one or more switches, including one or more interfaces for connecting to one or more external devices; a switch controller that executes software, including a routing algorithm; and a selectively configurable switch fabric connected to one or more memory sections and the one or more interfaces and interconnecting the memory sections and the one or more interfaces based on the routing algorithm; and a management system capable of receiving fault messages from the memory section controllers and inactivating the memory section corresponding to the fault message received by changing the routing algorithm, and wherein the management system is further capable of determining the routing algorithm for use by the selectively configurable switch fabric in interconnecting the memory sections and the one or more interfaces, and providing the routing algorithm to the switch controller.

144. Asserted claim 2 of the ’388 Patent recites a method for use in a storage system, comprising: storing data in storage locations in a memory device, the memory device included in a memory section; determining, by a management system, a routing algorithm for use by a switch controller that executes software, including the routing algorithm; providing, by the management system, the routing algorithm to the switch controller; executing, by the switch controller, the routing

algorithm, to configure a configurable switch connecting the memory section to an interface; detecting a fault associated with the data in the storage locations in the memory device; determining, by the management system in response to the detecting, a new routing algorithm that redirects data for the memory device to a replacement memory device; and providing the new routing algorithm to the switch controller.

145. In normal operation of the Accused Instrumentalities, the management system determines a new routing algorithm that redirects data for the memory device to a replacement memory device in response to detecting a fault.

146. During normal operation and upon detecting a dead DataNode (e.g., a DataNode with no heartbeat) the NameNode daemon (management system) bypasses the dead DataNode and sends I/O requests to other DataNodes storing replicas of blocks that were stored on the dead DataNode. The NameNode then schedules creation of new block replicas (to be stored on replacement memory devices) which result in an updated HDFS NameSpace (new routing algorithm).

147. Upon detecting a corrupted block (via DataBlockScanner) the NameNode daemon (management system) marks the block replica as corrupt and then schedules a copy of the block to be replicated (stored on replacement memory devices) on another datanode, so its replication factor is back at the expected level.

This results in an updated HDFS NameSpace (new routing algorithm).

148. During normal operation, the DataBlockScanner object creates a list of replicas that serves as the initial list of data blocks that it will scan for errors. When the NameNode becomes aware that a block is corrupt, it updates its internal tables to indicate that a block on a specific DataNode is corrupt and enters the corrupt replica into a list of blocks needing additional replicas. Once the replica has been created, the identity of the new replicas in this DataNode are sent to the NameNode.

149. When the NameNode daemon detects a fault (e.g. a dead NameNode or corrupt data block) an updating of the HDFS NameSpace is triggered that results in updates to the NameNode NameSpace (a new routing algorithm provided to the switch controller).

150. E*Trade is on notice of the infringing products, services, features, and how E*Trade operates the Accused Instrumentalities to perform the claimed methods and use the claimed apparatuses.

151. E*Trade's infringing conduct has damaged MNS.

152. E*Trade is liable to MNS in an amount that adequately compensates it for E*Trade's infringement, which, by law, can be no less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

NOTICE OF REQUIREMENT OF LITIGATION HOLD

153. E*Trade is hereby notified they are legally obligated to locate, preserve, and maintain all records, notes, drawings, documents, data, communications, materials, electronic recordings, audio/video/photographic recordings, and digital files, including edited and unedited or "raw" source material, and other information and tangible things that E*Trade knows, or reasonably should know, may be relevant to actual or potential claims, counterclaims, defenses, and/or damages by any party or potential party in this lawsuit, whether created or residing in hard copy form or in the form of electronically stored information (hereafter collectively referred to as "Potential Evidence").

154. As used above, the phrase “electronically stored information” includes without limitation: computer files (and file fragments), e-mail (both sent and received, whether internally or externally), information concerning e-mail (including but not limited to logs of e-mail history and usage, header information, and deleted but recoverable e-mails), text files (including drafts, revisions, and active or deleted word processing documents), instant messages, audio recordings and files, video footage and files, audio files, photographic footage and files, spreadsheets, databases, calendars, telephone logs, contact manager information, internet usage files, and all other information created, received, or maintained on any and all

electronic and/or digital forms, sources and media, including, without limitation, any and all hard disks, removable media, peripheral computer or electronic storage devices, laptop computers, mobile phones, personal data assistant devices, Blackberry devices, iPhones, video cameras and still cameras, and any and all other locations where electronic data is stored. These sources may also include any personal electronic, digital, and storage devices of any and all of Defendant's agents, resellers, or employees if E*Trade's electronically stored information resides there.

155. E*Trade is hereby further notified and forewarned that any alteration, destruction, negligent loss, or unavailability, by act or omission, of any Potential Evidence may result in damages or a legal presumption by the Court and/or jury that the Potential Evidence is not favorable to E*Trade's claims and/or defenses. To avoid such a result, E*Trade's preservation duties include, but are not limited to, the requirement that E*Trade immediately notify its agents and employees to halt and/or supervise the auto-delete functions of its electronic systems and refrain from deleting Potential Evidence, either manually or through a policy of periodic deletion.

PRAYER FOR RELIEF

MNS prays for the following relief:

a) A judgment be entered that E*Trade has directly and indirectly infringed one or more claims of the Asserted Patents;

b) A judgment be entered that the Asserted Patents are valid and enforceable;

c) An award of damages adequate to compensate MNS for E*Trade's infringement up until the date such judgment is entered, including prejudgment and post-judgment interest, costs, and disbursements as justified under 35 U.S.C. § 284 and an accounting, if necessary, to adequately compensate MNS for E*Trade's infringement;

d) A judgment that MNS be awarded attorneys' fees, costs, and expenses incurred in prosecuting this action; and

e) A judgment that MNS be awarded such further relief at law or in equity as the Court deems just and proper.

DEMAND FOR JURY TRIAL

MNS demands trial by jury for all issues so triable pursuant to Fed. R. Civ. P. 38(b) and Civil L.R. 3-6(a).

Dated: March 5, 2021

/s/Daniel A. Kent

Daniel A. Kent
dankent@kentrisley.com
Tel: (404) 585-4214
Fax: (404) 829-2412
Stephen R. Risley
steverisley@kentrisley.com
Tel: (404) 585-2101
Fax: (404) 389-9402
KENT & RISLEY LLC
5755 N Point Pkwy Ste 57
Alpharetta, GA 30022

Cabrach J. Connor
(to be admitted pro hac vice)
cab@connorkudlaclee.com
Kevin S. Kudlac
Georgia Bar No. 429946
kevin@connorkudlaclee.com
CONNOR KUDLAC LEE PLLC
609 Castle Ridge Road, Suite 450
Austin, TX 78746
Telephone: 512.777.1254
Facsimile: 888.387.1134

Attorneys for Plaintiff
Mobile Networking Solutions, LLC